

# Search for a Higgs Portal scalar decaying to electron-positron pairs in MicroBooNE

## MICROBOONE-NOTE-1092-PUB

The MicroBooNE Collaboration<sup>a</sup>  
July 2020

We present a search for the decays of a neutral scalar boson produced by kaons decaying at rest, in the context of the Higgs Portal model, using the MicroBooNE detector. We analyze data triggered in time with the Fermilab NuMI neutrino beam spill, with an exposure of  $1.93 \times 10^{20}$  protons on target. We look for monoenergetic scalars coming from the direction of the NuMI hadron absorber, 100 m away from the detector, and decaying to electron-positron pairs. We observe 5 candidate events, with a Standard Model background prediction of  $2.0 \pm 0.8$ . We set an upper limit on the scalar–Higgs mixing angle  $\theta < (4.3 - 5.8) \times 10^{-4}$  at the 95% confidence level, for scalar masses in the range (100 – 200) MeV/ $c^2$ . We exclude at the 95% confidence level the remaining model parameters required to explain the central value of the KOTO anomalous excess of  $K_L^0 \rightarrow \pi^0 +$  invisible decays using this model.

### I. INTRODUCTION

The Higgs Portal model [1] is an extension to the Standard Model, where an electrically-neutral real singlet scalar boson ( $S$ ) mixes with the Higgs boson with mixing angle  $\theta$ . Through this mixing, it acquires a coupling to Standard Model fermions via their Yukawa couplings with the Higgs boson, and proportional to  $\sin \theta$ . For the scalar mass in the range (100 – 200) MeV/ $c^2$ , and assuming that there are no new dark sector particles lighter than half its mass,  $S$  will decay to electron-positron pairs with partial width [2]

$$\Gamma = \theta^2 \frac{m_e^2 m_S}{8\pi v^2} \left(1 - \frac{4m_e^2}{m_S^2}\right)^{\frac{3}{2}}, \quad (1)$$

where  $m_S$  is the scalar mass,  $m_e$  the electron mass, and  $v$  the Higgs field vacuum expectation value. For these masses,  $S$  can be produced from kaon two-body decays in association with pions, with the dominant production process being a penguin diagram with a top quark running in the loop. The partial width of the production process is [2]

$$\Gamma \simeq \frac{\theta^2}{16\pi m_K} \left| \frac{3V_{td}^* V_{ts} m_t^2 m_K^2}{32\pi^2 v^3} \right|^2 \lambda^{1/2} \left(1, \frac{m_S^2}{m_K^2}, \frac{m_\pi^2}{m_K^2}\right), \quad (2)$$

where  $m_K$  is the kaon mass,  $m_\pi$  the pion mass,  $m_t$  the top quark mass,  $V_{td}$  and  $V_{ts}$  the elements of the CKM matrix, and  $\lambda$  the Källén Lambda function.

Recently, the KOTO collaboration has reported [3] the anomalous excess of  $K_L^0 \rightarrow \pi^0 +$  invisible decays, two orders of magnitude more frequent than the Standard Model prediction for  $K_L^0 \rightarrow \pi^0 \nu \bar{\nu}$  decays. The Higgs Portal model could explain the high rate of these decays, with the required model parameter  $\theta$  value  $\sim (5 - 7) \times 10^{-4}$  over the scalar mass range (100 – 200) MeV/ $c^2$  to agree with KOTO central value branching ratio [4]. There are experimental limits that exclude this central value for  $m_S < 118$  MeV/ $c^2$  by the E949 experiment [5] and  $m_S > 160$  MeV/ $c^2$  by the NA62 experiment [6]. Data from the CHARM [7] experiment is also sensitive to this model, however this requires reinterpretation of the experiment’s search for axion-like particles decaying to electron positron pairs, and different phenomenology groups [4, 8, 9] have different estimates of the kaon fluxes and scalar acceptances observed by CHARM. Taking the most sensitive of these estimates [4], the KOTO central value can be excluded for  $m_S > 150$  MeV/ $c^2$ .

This note presents the first search for Beyond the Standard Model electron-positron pair production in a liquid argon time projection chamber, using the MicroBooNE detector. We can use this search to exclude at 95% confidence the remaining parameter space of the Higgs Portal model required to explain the central value of the KOTO excess.

### II. EXPERIMENTAL SETUP

The MicroBooNE experiment is primarily designed for neutrino scattering measurements in Fermilab’s *Booster Neutrino Beam* (BNB). The detector sits just below surface, and comprises an 85 ton liquid argon time projection

---

<sup>a</sup> email: MICROBOONE.INFO@fnal.gov

chamber (TPC) with active dimensions of 2.6 m along the drift direction (horizontal and perpendicular to the beam axis), 2.3 m in the vertical direction, and 10.4 m along the direction parallel to the BNB direction. Three wire planes with 3 mm pitch and oriented at  $60^\circ$  rotations relative to each other, record signals of drifted ionization electrons. An array of 32 eight inch photomultiplier tubes (PMTs) distributed behind the wire planes provides timing information for scintillation signals inside the TPC. Part-way through the detector operations, a cosmic ray tagger (CRT) system was installed, with four walls of plastic scintillator panels situated along the top, bottom, and long sides of the cryostat providing timing coincidence signals for some cosmic rays entering the TPC. Further details on the detector are given in Refs. [10] and [11].

In addition to being on the BNB beamline, the MicroBooNE detector is also situated close to Fermilab’s *Neutrinos at the Main Injector* (NuMI) neutrino beam [12]. The detector is  $8^\circ$  off-axis from the NuMI target, located 680 m away. The Main Injector delivers 120 GeV protons that hit the graphite target, producing secondary hadrons. A system of electromagnetic horns focus positive and negative mesons either towards or away from the beam axis, depending on the horn polarity. In Forward Horn Current (FHC) mode, positive mesons are bent towards the axis, producing a beam of mostly neutrinos from the meson decays, whereas in Reverse Horn Current (RHC) mode, the negative mesons are focused towards the beam axis, and a beam of mostly antineutrinos is produced by the meson decays. A 675 m long helium-filled decay volume is situated downstream of the target and horn system, at the end of which is a 5 m deep hadronic absorber. Any surviving mesons will be stopped in the absorber, and will be at rest if they decay. Protons that didn’t interact in the target will also travel down the decay volume and produce more secondary mesons in the absorber. Of any kaons produced, only  $K^+$  will decay at rest, with  $K^-$  generally absorbed on nuclei, and  $K^0$  reinteracting hadronically or decaying in flight. The absorber is at a distance of 100 m from the MicroBooNE detector, at an angle of  $\sim 55^\circ$  with respect to the BNB direction, such that any particles entering MicroBooNE from the absorber are entering in a backward direction compared to most neutrino interactions seen by the detector.

We exploit the unique signature of scalars produced by charged kaons decaying at rest (KDAR) in the NuMI hadron absorber to search for evidence of the Higgs Portal scalar model. The scalars will be produced monoenergetically in the two-body decay, and have a kinetic energy  $\sim 100$  MeV, dependent on the scalar mass. Due to the time of flight of the hadron beam to the absorber, followed by the time of flight of the scalars into the detector, the scalar decay signal is expected to be delayed with respect to the neutrino backgrounds by around 600 ns. Once they decay, the scalars produce an electron-positron pair with  $\sim 250$  MeV summed energy, and a strongly peaked opening angle distributions (peaking from  $\sim 50^\circ$  to  $\sim 100^\circ$  depending on the scalar mass, in the mass range considered) due to the monoenergetic boost of the isotropic back-to-back decay in the rest frame. The selection in this analysis will rely on searching for the characteristic kink between two reconstructed electron candidates, pointing in a direction consistent with being from the NuMI hadron absorber.

### III. DATA, SIMULATION, AND RECONSTRUCTION

We analyse only a small fraction of the total available dataset. We use  $0.92 \times 10^{20}$  protons on target (POT) of exposure during Run 1 of MicroBooNE’s operations (during 2015–2016), and  $1.01 \times 10^{20}$  POT of Run 3 data (2017–2018), out of a total  $22.3 \times 10^{20}$  POT of NuMI data collected by MicroBooNE over its lifetime. During the Run 1 dataset period, the NuMI beam operated in FHC mode, and during the Run 3 period, RHC mode was used. The CRT had been fully installed by the Run 3 data taking period. The On-beam data is read out from the detector (an ‘event’) when there is a NuMI beam spill timing signal sent by the Fermilab accelerator complex. An on-line trigger is employed to save to tape only those events which pass optical trigger criteria, requiring at least one PMT to fire in time with the beam, with a total integrated charge, summed over all PMTs in a 100 ns window, above a configurable photoelectron threshold.

To simulate the signal, we use the `g4numi` program [13] which employs a GEANT4 [14] simulation of the NuMI beamline. We use this to produce the position and timing distribution of kaons decaying at rest in the NuMI hadron absorber. For the rate of charged kaon decays at rest, we use the equivalent of the MiniBooNE estimate [15] of 0.085 KDAR- $\nu_{\mu,s}$  produced in the NuMI hadron absorber per POT, with a 30% uncertainty on this value. The scalars are emitted isotropically from the kaon decay positions, and the Lorentz-boosted lifetime, calculated from the scalar’s total width, is used to correctly distribute the scalar decay position, keeping only those that decay within the detector active volume. The scalars decay to back-to-back electrons and positrons isotropically in the rest frame, and are boosted by the scalar’s momentum.

The `g4numi` program is used to simulate the flux of background neutrinos intersecting the detector, and we use PPFX [13, 16] to correct the central value flux prediction and provide flux uncertainties. The GENIE [17] software version 3.0.4, with a MicroBooNE tune of the cross-section models [18], is used to calculate the neutrino interaction cross-sections and final state kinematics. There are two types of neutrino background simulations, one with neutrino interactions located within the detector cryostat, and a simulation of neutrinos interacting outside the detector

(commonly referred to as ‘dirt’ interactions), with the secondary products entering the detector from outside.

For both signal and background simulations, the decay or interaction products are propagated through a GEANT4 simulation of the detector, and the response of the detector to both light and charge is simulated.

To estimate the cosmic-induced backgrounds, we have a dataset of events produced out of time with the beam, but employing the same trigger thresholds, called the Off-beam dataset. In addition there is an unbiased dataset of out-of-time events but without trigger conditions. The unbiased dataset is used for the simulation, whereby the simulated detector response to neutrino interactions and scalar decays is overlaid on top of the cosmic background in the unbiased event, decalibrating the simulation according to the underlying unbiased event and applying equivalent detector effects such as dead channels. A trigger emulator is applied to the summed optical information.

All three types of datasets (On-beam, Off-beam, and simulated) are propagated through a chain of reconstruction algorithms. The optical reconstruction uses the PMT waveforms to produce ‘flashes’ of coincident PMT hits. For the TPC information, there is a two-dimensional deconvolution of the signal waveforms on the wires within each plane [19]. Hits are formed from a Gaussian peak finding algorithm applied along the wire waveform. The Pandora framework [20] uses particle flow algorithms to cluster the hits of a single plane, and then match clusters across planes into three-dimensional reconstructed objects. Pandora uses an independently trained multivariate classifier to produce a track/shower score in the range  $[0, 1]$  for each object, and treats objects with this score  $< 0.5$  as showers, and  $> 0.5$  as tracks. Pandora also ‘slices’ the event up into groups of reconstructed objects that it considers to be independent interactions (either cosmogenic or beam-induced), and removes obvious cosmic slices. For remaining slices, there is a flash matching algorithm that produces a PMT hit hypothesis using the reconstructed objects in the slice, and attempts to match it with the observed flashes in the beam timing window, producing a  $\chi^2$  value for the best match. The best matching slice, if any, is labelled as the *neutrino slice*.

#### IV. SIGNAL SELECTION

The selection was developed blindly, using only the simulated and the Off-beam cosmic background datasets. There is a two-step approach to the selection, first to pre-select decay candidates (pairs of reconstructed objects), and then to train multiple boosted decision trees (BDTs) to filter for signal-like decay candidates using the geometric information of the objects, exploiting the characteristic kink topology of signal decays.

The preselection uses neutrino slices (as labelled by Pandora) in the event. The slice has to be matched to a PMT flash with a time of  $[5.8, 16.8] \mu\text{s}$  (where the NuMI prompt neutrino spill produces flashes in the range  $[6.1, 15.7] \mu\text{s}$ ), and the flash-matching  $\chi^2$  has to be less than 10. For the data taking during Run 3, after the CRT was installed, events cannot have a CRT hit in coincidence with the beam timing. The total number of objects in the slice has to be  $\leq 5$ , and of these, a maximum of four can be labelled as tracks. For all possible pairs of objects in the slice, the minimum distance between the object vertices (for reconstructed tracks, both start or end positions, and for showers, only start positions) is calculated, and if this distance is less than 5 cm, a decay vertex is produced at the mid-point between the object vertices with the minimum separation. This reconstructed decay vertex has to be within the active volume of the detector. Slices with more than two object could conceivably form multiple decay candidates, all of which are preselected and passed through the BDT selection.

We apply two different BDTs to the preselected candidates: one trained against the cosmic backgrounds, and one trained against the neutrino interactions simulated in the cryostat. Each of these is trained separately over the Run 1 events and Run 3 events, meaning that there are four BDTs in total. We split up the two run periods, because the use of the CRT, and the differences between neutrino and antineutrino beam running, can change the preselected topologies and other background distributions that the BDTs are trained against. We use `xgboost` [21] version 1.3 to train and apply the BDTs. We choose to train the BDTs on ten input distributions each, which are chosen as the ten most important inputs after a preliminary round of training with 81 inputs. Nine of the ten input values are the same for both the cosmic-focused and neutrino-focused BDTs. These are: the opening angle between the two reconstructed objects; the opening angle in the plane transverse to the hadronic absorber direction; the two angles between the two objects and the hadronic absorber direction; the Pandora track/shower score of the larger of the two objects by number of hits; the number of hits of that larger object by number of hits; the total number of hits contained in other objects in the slice, not including the two objects forming the decay candidate; the maximum height of shower start positions or track start or end positions, for any other objects in the slice, relative to the decay vertex (‘*y*-extent’); and the minimum longitudinal position of shower start positions or track start or end positions, for any other objects in the slice, relative to the decay vertex (‘*z*-extent’). The last two variables can be marked as ‘missing’ in the case that the slice contains only two objects, and `xgboost` can cope with missing values. The cosmic-focused BDT has as its tenth input variable, the length of the larger object of the two objects by number of hits. The neutrino-focused BDT has as its tenth input, the number of tracks in the slice.

The BDTs are trained on a specially generated signal simulation, where each decay is of a random scalar mass

Uncertainty	Background	Signal
Flux (hadron production)	34.0%	30.0%
Cross-section model	3.2%	–
Off-beam statistics	67.4%	–
Simulated statistics	24.3%	< 2.2%

TABLE I. Relative sizes of systematic uncertainties for the signal and background model, in the signal region defined by the optimal BDT cuts. The signal simulation statistical error varies depending on the mass point.

uniformly chosen between 100 and 200 MeV/ $c^2$ , in order to make the BDTs approximately mass-agnostic. The training decay candidates have to be well-reconstructed, with the cosmic contamination of both objects below 10% each, and the reconstructed vertex and directions close to the truth. The neutrino-focused BDT is trained against 10% of the simulated cryostat neutrino statistics, with the other 90% along with all the dirt simulation used for the sensitivity and limit calculations; the candidate reconstructed objects for the neutrino background training sample have to have cosmic contamination below 10% each, similarly to the signal sample. The cosmic-focused BDT is trained against Off-beam candidates failing the flash-matching  $\chi^2$  cut. In all three cases, the training datasets are divided up in a 7 : 3 ratio as training and testing samples.

We have evaluated the performance of the simulation with respect to the input variables by examining their distributions in a control region where no signal is expected. Because the signal decays are expected to be delayed with respect to the neutrino beam, the control region is events with an early flash time, in the range [6.1, 6.8]  $\mu\text{s}$  for Run 1, and [6.1, 6.9]  $\mu\text{s}$  in Run 3. However, because we do not know the fraction of POT delivered in this early time window (only that delivered over the total beam spill), we cannot check the normalisation, only the shape. We see good agreement across all variables. We also look at the performance of the BDT score distributions in this control region, and see good agreement as well.

We consider several sources of systematic uncertainty. The flux normalisation uncertainty on the signal model is set to 30% as used by MiniBooNE [15]. The uncertainty on the background neutrino flux and cross-section modelling is evaluated by reweighting events, whereby hadron production uncertainties (for the flux), and physics model parameters (for the cross-section), were varied 100 times and 441 times respectively, and a weight calculated for each simulated event between the central value and the modified model. We use PPFX for the flux reweighting, and tools included with GENIE [22] for the cross-section reweighting. The uncertainties on observable distributions or event counts is calculated from the standard deviation on the observable across the 100 or 441 variations. We also include the Monte Carlo statistical error and Off-beam data statistical error, as uncertainties for the model prediction. We have not yet evaluated detector systematic uncertainties, as that requires a time-intensive data processing campaign to re-simulate or re-reconstruct events. However, the final search is statistics-dominated, that even a detector systematic uncertainty as large as the other ones summed in quadrature, and anti-correlated between signal and background, does not change the sensitivity or observed limit by a significant amount, and does not change our conclusions.

We choose the optimal BDT cuts to maximise the sensitivity of the selection to the model parameter  $\theta$  for 100 MeV/ $c^2$  mass scalar bosons, as we expect even better sensitivity at higher masses. We perform a single-bin counting experiment, summing the event yields from both Runs after BDT cuts, and using correlated flux and cross-section uncertainties. The 95% confidence level (CL) sensitivity and limit is calculated with the modified-frequentist CLs method, using the RooStats statistical package [23], including systematic uncertainties as constrained Gaussian nuisance terms. We scan over the four-dimensional BDT cut space to find the cut values which maximise the sensitivity. The uncertainties in the signal region after the optimal cuts are given in Table I. The final sensitivity we obtain is presented in Figure 1.

## V. RESULTS AND CONCLUSION

After unblinding, we observe six events in the signal region, one of which we reject because its flash time is in the 0.3  $\mu\text{s}$  window before the start of the beam spill, and so is an obvious cosmic background interaction. After hand-scanning the other five events, we observe that none of them appear to be obvious electron-positron pairs, and they are all consistent with neutrino or cosmic background interactions. However we do not apply any additional post-unblinding selection to remove these, as that would require a re-evaluation of efficiencies and systematic uncertainties due to hand scans.

The five observed events are consistent with the background prediction of  $2.0 \pm 0.8$ . The 95% CL observed upper limit of this search to the Higgs Portal model is presented in Figure 1, along with the median and  $\pm 1\sigma$  expected sensitivity. The upper limit is presented in the context of the parameters required to explain the KOTO anomalous

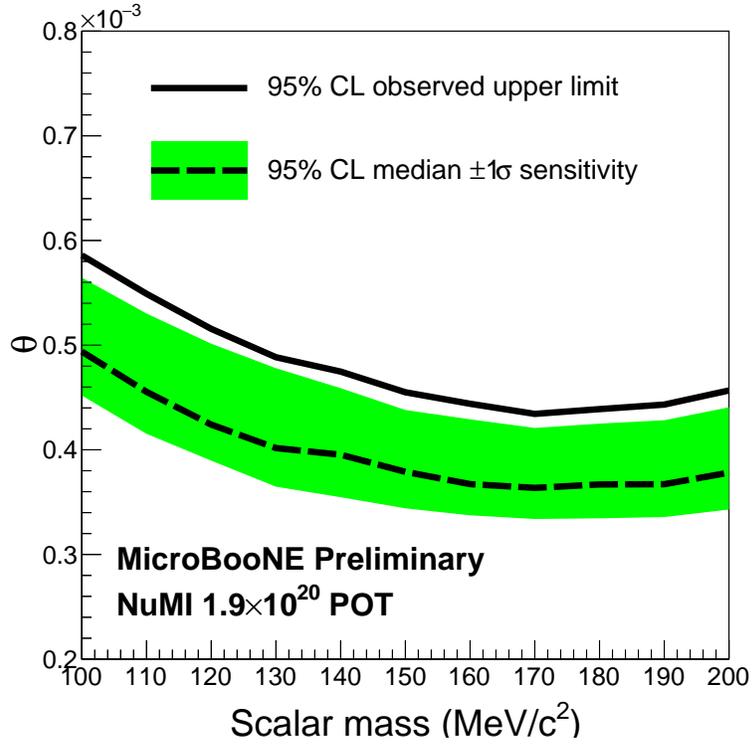


FIG. 1. 95% confidence level sensitivity and observed limit of this search to the Higgs Portal model parameter  $\theta$  and the scalar mass.

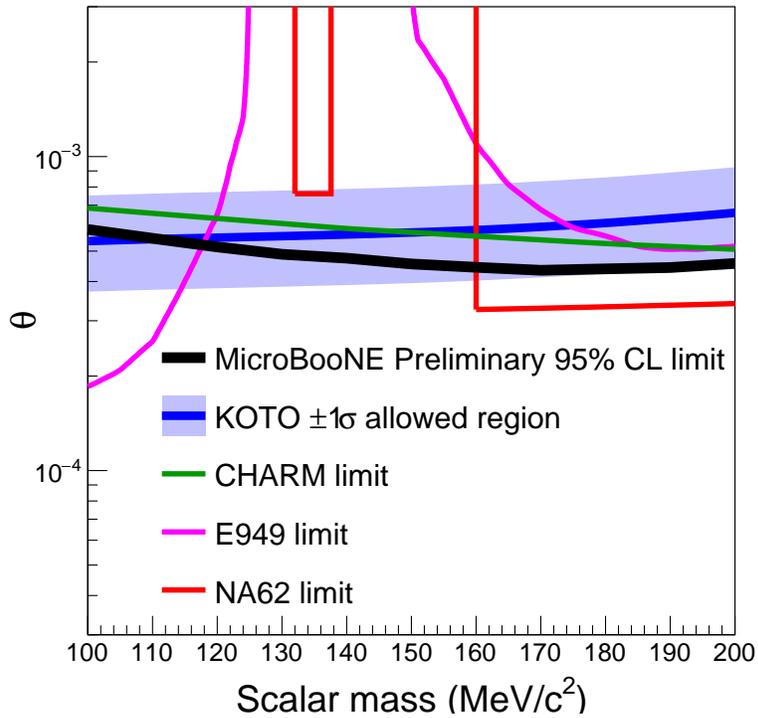


FIG. 2. This result in the context of the KOTO central value [3], and exclusions from E949 [5], NA62 [6], and CHARM [7]. Figure adapted from Ref. [4].

excess, along with other experimental limits, in Figure 2. The limits obtained by MicroBooNE can rule out the whole remaining parameter space for the KOTO central value, at the 95% confidence level.

We plan to analyse the full collected dataset, including expanding the search for scalars produced by kaons decaying in flight in both the NuMI and BNB beamlines. Combining all the datasets, MicroBooNE should have the reach [2] to exclude the whole  $2\sigma$ -allowed parameter space to explain the KOTO anomaly with the Higgs Portal model. We will also search for heavier scalar masses, above the di-muon threshold. For future iterations of this KDAR analysis, we plan to develop the preselection to reject obvious neutrino interactions from the outset. We hope to better understand the types of background that populate the signal region, so that an efficient targeted simulation can be developed to generate more statistics in the tails of the distributions leaking into the signal region.

## VI. ACKNOWLEDGEMENTS

We thank Brian Batell, Joshua Berger, and Ahmed Ismail, for providing the motivation and outlining the KDAR search strategy, in a recent phenomenological study [2].

## VII. SUPPLEMENTARY MATERIAL

An example simulated display of a  $150\text{MeV}/c^2$  scalar decaying inside MicroBooNE is shown in Figure 3. Plots of the BDT input variable distributions for the preselected decay candidates are shown in Figures 4 to 14. The BDT score distributions are presented in Figures 15 and 16.

There is a deficit of the On-beam data with respect to the model prediction, however this deficit is within the flux uncertainty error bands. If this deficit is interpreted as a flat inefficiency on both signal and background, the observed limit on  $\theta$  is weakened by  $\sim 2 \times 10^{-5}$ , and doesn't change the main conclusions of this work.

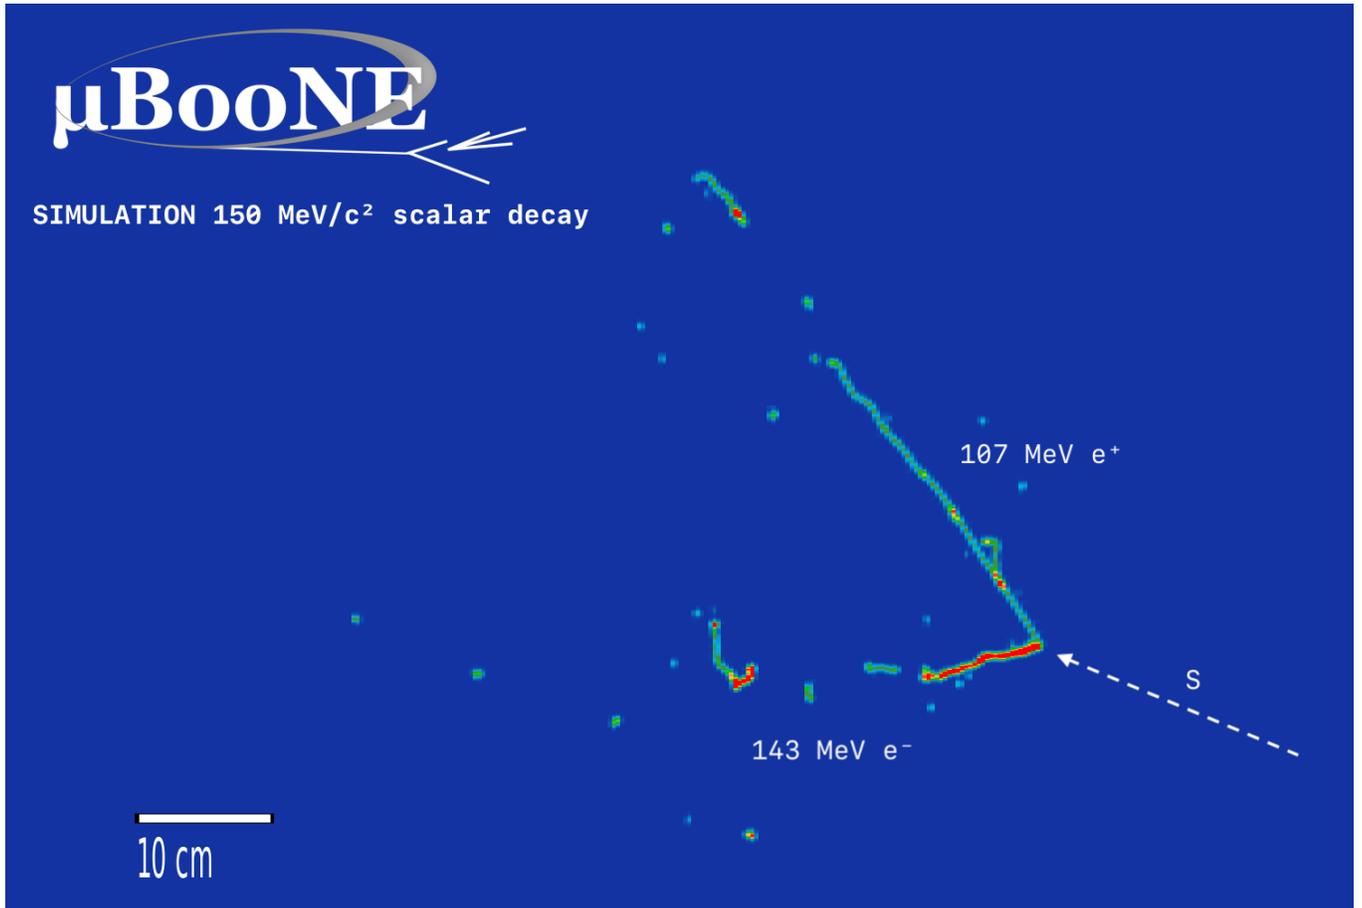


FIG. 3. Event display of a simulated  $150\text{ MeV}/c^2$  scalar coming from the NuMI hadron absorber, decaying to an electron-positron pair.

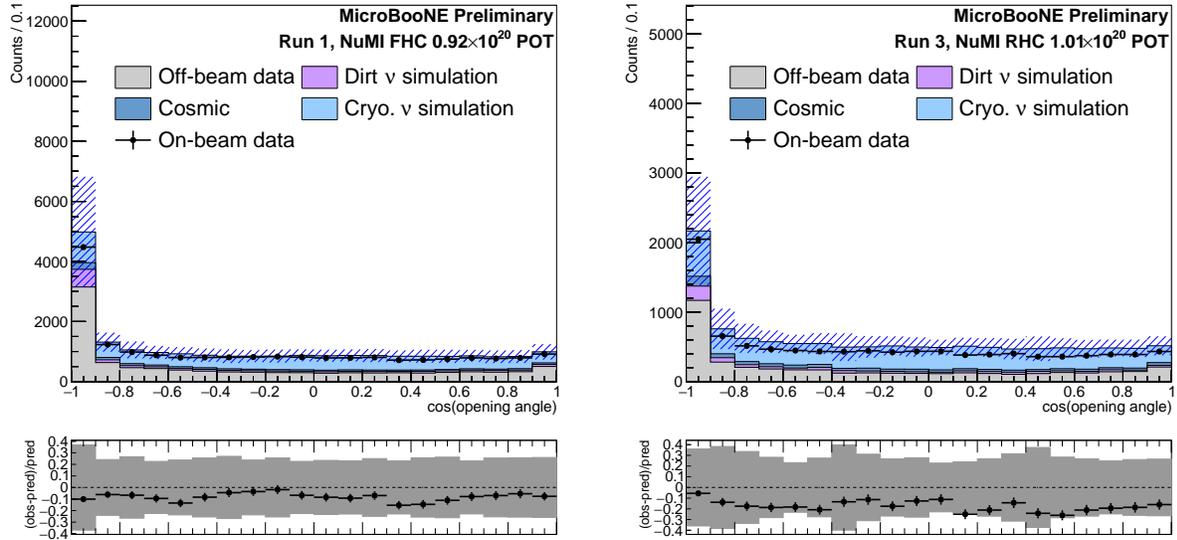


FIG. 4. Distribution showing the opening angle of the two objects forming the decay candidate for Run 1 (left) and Run 3 (right). The ‘cosmic’ background is cosmic contamination from the underlying event in simulated data. The hashed area is the systematic error band due to uncertainties in the flux and cross-section models, and simulation statistics. The ratio plot shows the total uncertainty in grey. A 10% (15%) deficit of data with respect to the model prediction in Run 1 (Run 3) is within the flux modelling uncertainty.

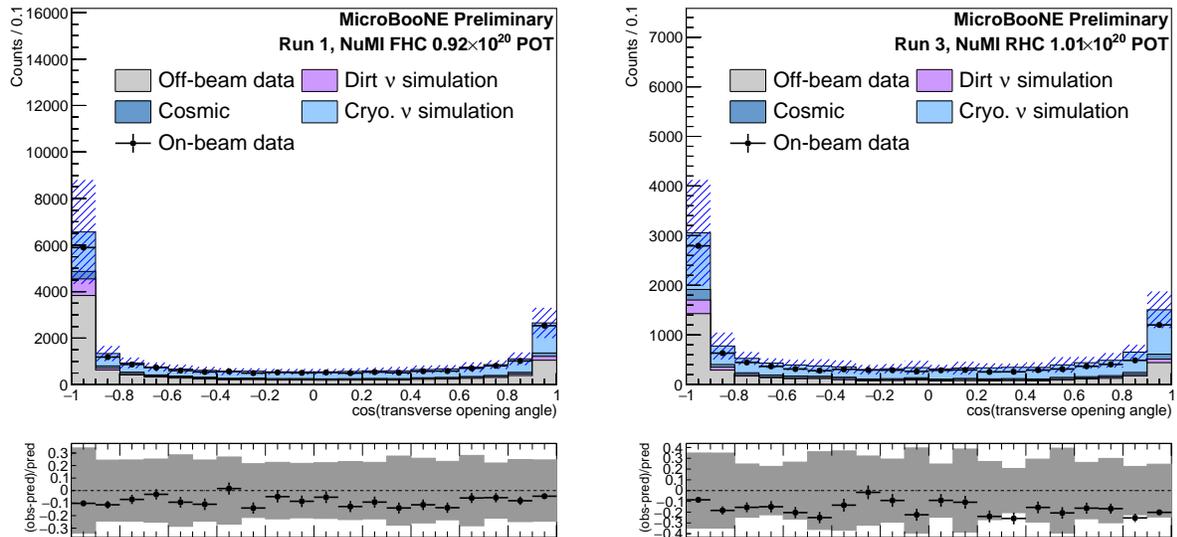


FIG. 5. Distribution showing the opening angle in the plane transverse to the assumed scalar direction for Run 1 (left) and Run 3 (right). The ‘cosmic’ background is cosmic contamination from the underlying event in simulated data. The hashed area is the systematic error band due to uncertainties in the flux and cross-section models, and simulation statistics. The ratio plot shows the total uncertainty in grey. A 10% (15%) deficit of data with respect to the model prediction in Run 1 (Run 3) is within the flux modelling uncertainty.

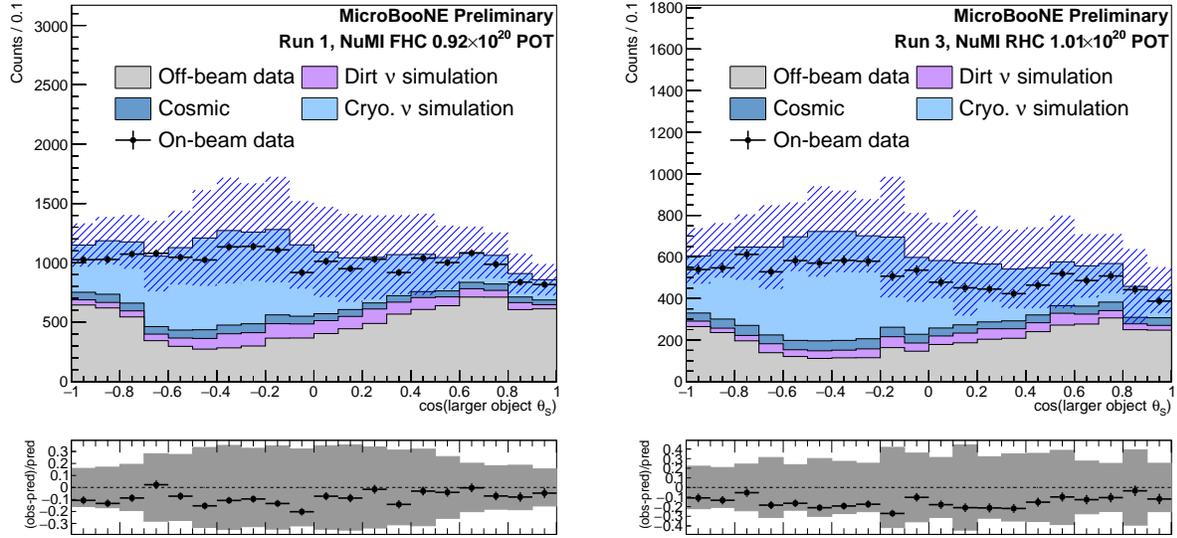


FIG. 6. Distribution showing the angle of the larger object (by number of hits) with respect to the assumed scalar direction for Run 1 (left) and Run 3 (right). The ‘cosmic’ background is cosmic contamination from the underlying event in simulated data. The hashed area is the systematic error band due to uncertainties in the flux and cross-section models, and simulation statistics. The ratio plot shows the total uncertainty in grey. A 10% (15%) deficit of data with respect to the model prediction in Run 1 (Run 3) is within the flux modelling uncertainty.

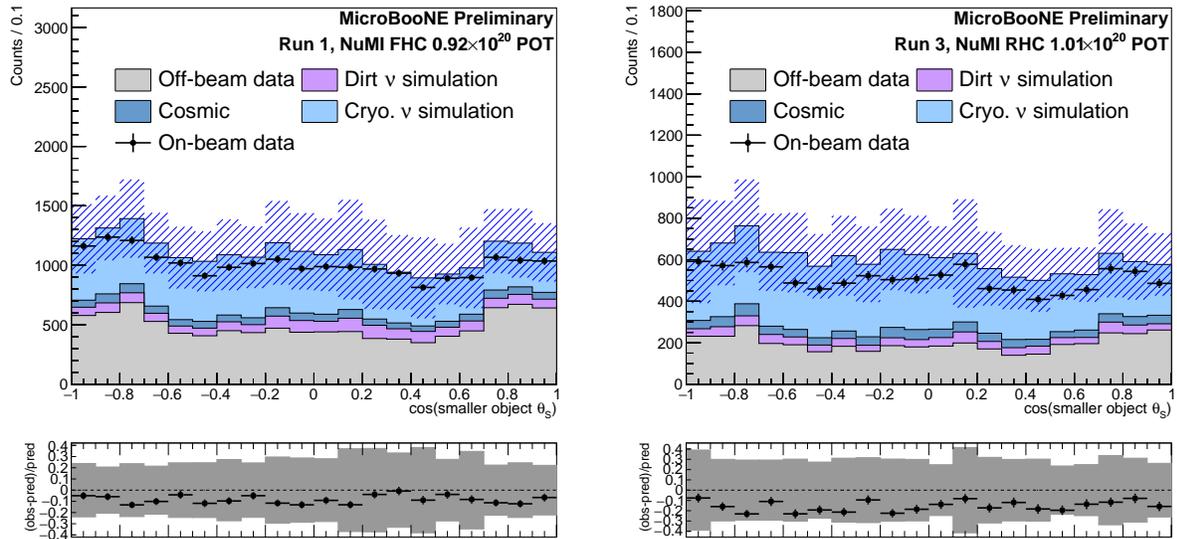


FIG. 7. Distribution showing the angle of the smaller object (by number of hits) with respect to the assumed scalar direction for Run 1 (left) and Run 3 (right). The ‘cosmic’ background is cosmic contamination from the underlying event in simulated data. The hashed area is the systematic error band due to uncertainties in the flux and cross-section models, and simulation statistics. The ratio plot shows the total uncertainty in grey. A 10% (15%) deficit of data with respect to the model prediction in Run 1 (Run 3) is within the flux modelling uncertainty.

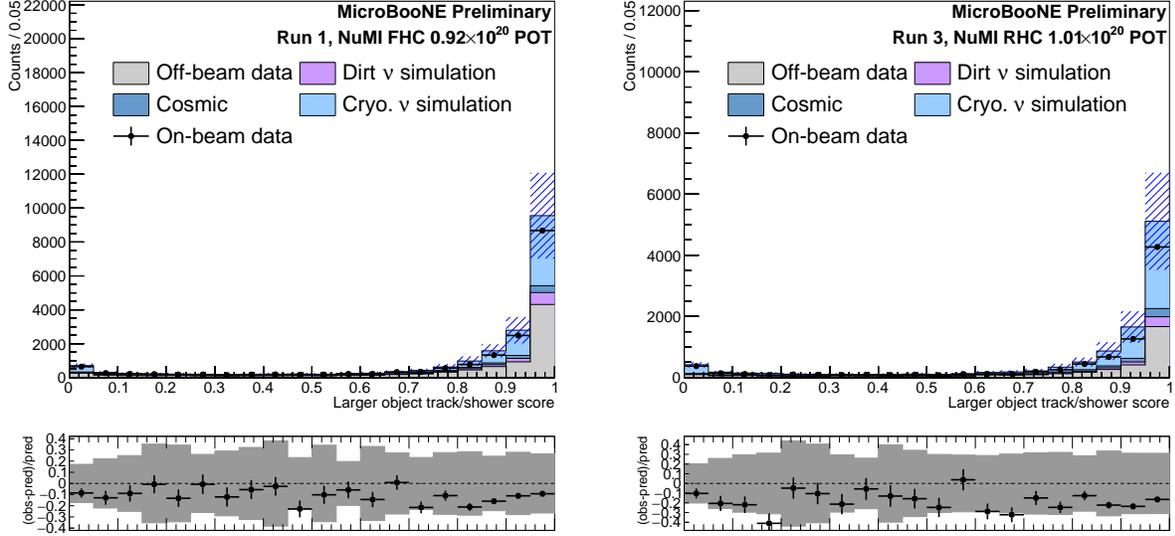


FIG. 8. Distribution showing the track/shower score of the larger object (by number of hits) for Run 1 (left) and Run 3 (right). The ‘cosmic’ background is cosmic contamination from the underlying event in simulated data. The hashed area is the systematic error band due to uncertainties in the flux and cross-section models, and simulation statistics. The ratio plot shows the total uncertainty in grey. A 10% (15%) deficit of data with respect to the model prediction in Run 1 (Run 3) is within the flux modelling uncertainty.

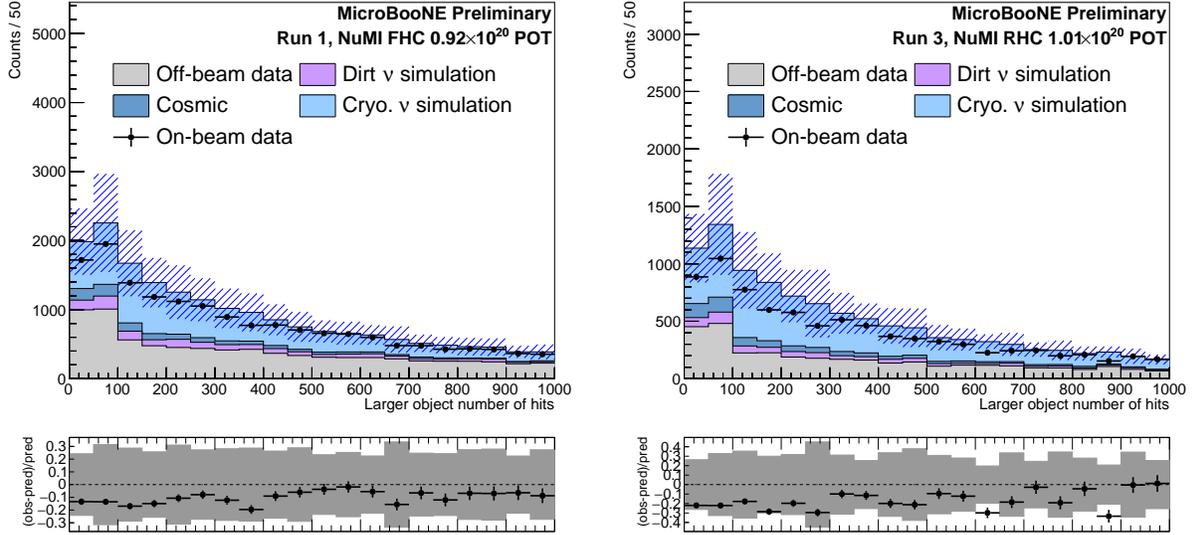


FIG. 9. Distribution showing the number of hits in the larger object (by number of hits) for Run 1 (left) and Run 3 (right). The ‘cosmic’ background is cosmic contamination from the underlying event in simulated data. The hashed area is the systematic error band due to uncertainties in the flux and cross-section models, and simulation statistics. The ratio plot shows the total uncertainty in grey. A 10% (15%) deficit of data with respect to the model prediction in Run 1 (Run 3) is within the flux modelling uncertainty.

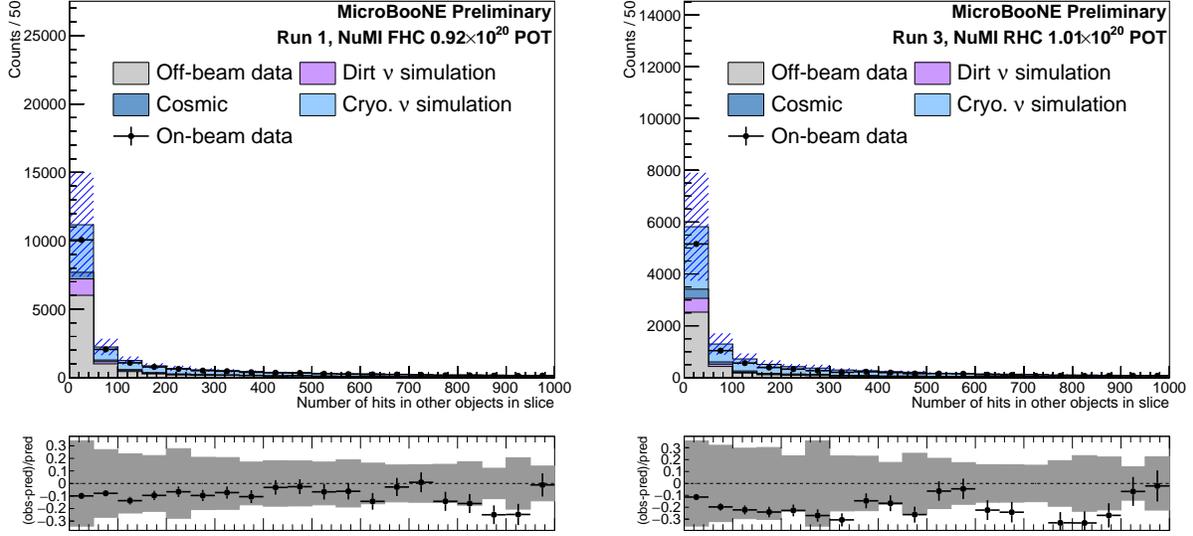


FIG. 10. Distribution showing the number of hits in other objects of the slice for Run 1 (left) and Run 3 (right). The ‘cosmic’ background is cosmic contamination from the underlying event in simulated data. The hashed area is the systematic error band due to uncertainties in the flux and cross-section models, and simulation statistics. The ratio plot shows the total uncertainty in grey. A 10% (15%) deficit of data with respect to the model prediction in Run 1 (Run 3) is within the flux modelling uncertainty.

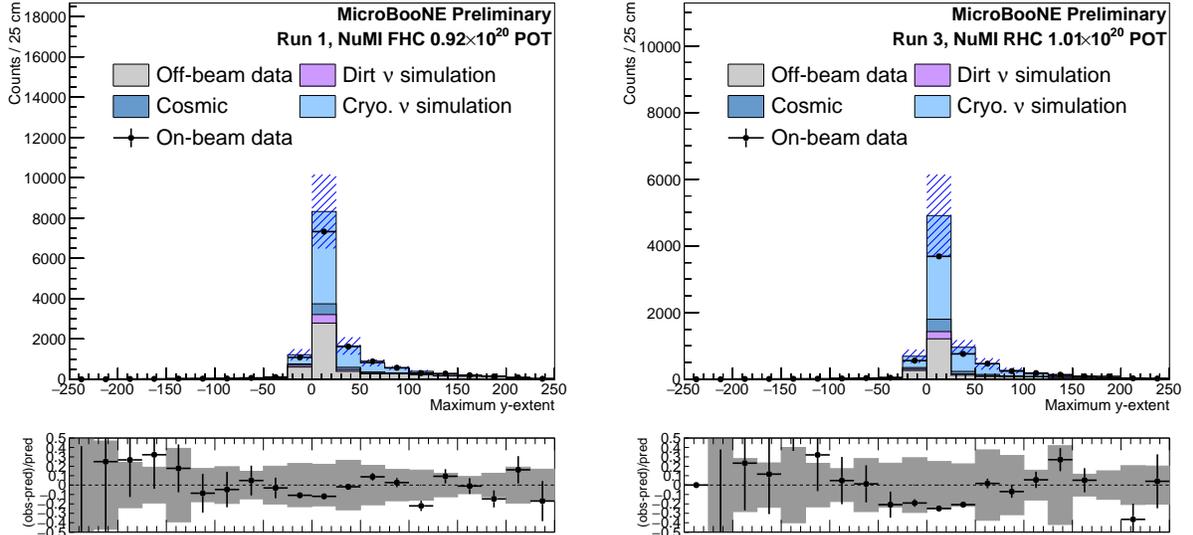


FIG. 11. Distribution showing the maximum  $y$ -extent for Run 1 (left) and Run 3 (right). The ‘cosmic’ background is cosmic contamination from the underlying event in simulated data. The hashed area is the systematic error band due to uncertainties in the flux and cross-section models, and simulation statistics. The ratio plot shows the total uncertainty in grey. A 10% (15%) deficit of data with respect to the model prediction in Run 1 (Run 3) is within the flux modelling uncertainty.

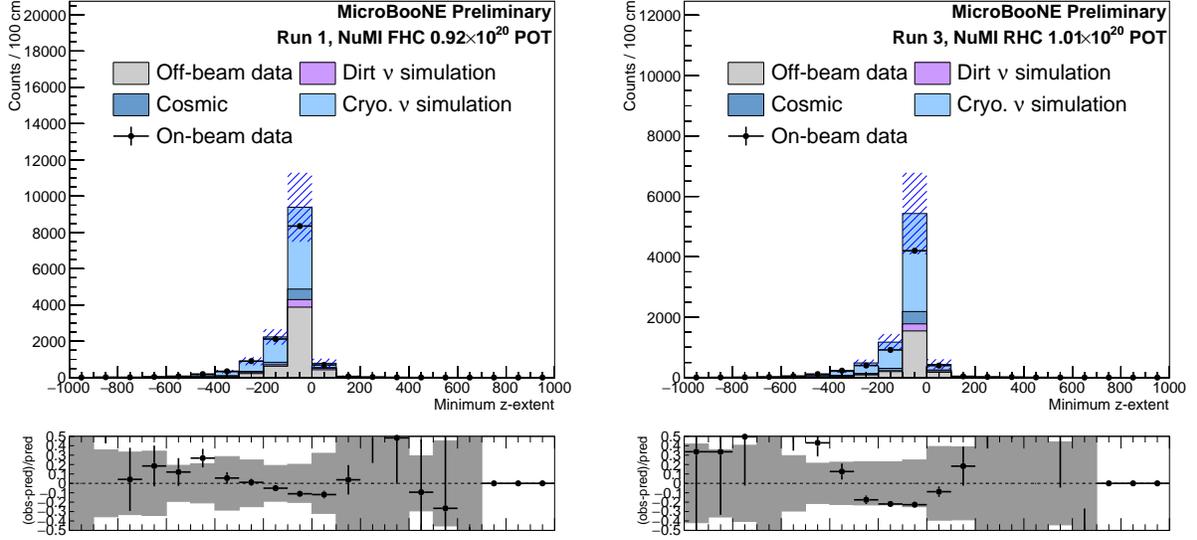


FIG. 12. Distribution showing the minimum  $z$ -extent for Run 1 (left) and Run 3 (right). The ‘cosmic’ background is cosmic contamination from the underlying event in simulated data. The hashed area is the systematic error band due to uncertainties in the flux and cross-section models, and simulation statistics. The ratio plot shows the total uncertainty in grey. A 10% (15%) deficit of data with respect to the model prediction in Run 1 (Run 3) is within the flux modelling uncertainty.

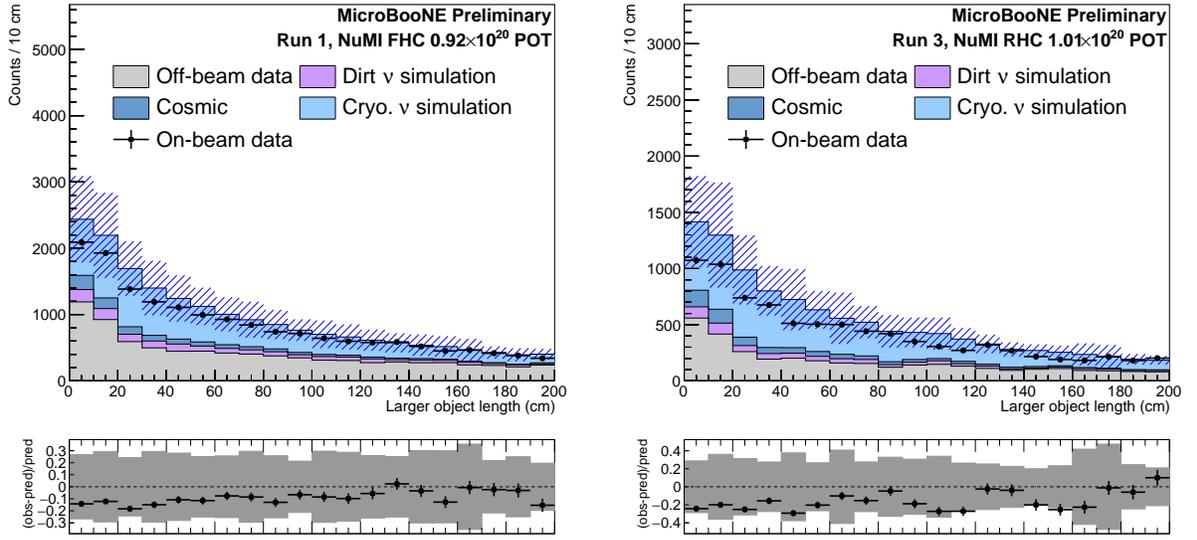


FIG. 13. Distribution showing the length of the larger object (by number of hits) for Run 1 (left) and Run 3 (right). The ‘cosmic’ background is cosmic contamination from the underlying event in simulated data. The hashed area is the systematic error band due to uncertainties in the flux and cross-section models, and simulation statistics. The ratio plot shows the total uncertainty in grey. A 10% (15%) deficit of data with respect to the model prediction in Run 1 (Run 3) is within the flux modelling uncertainty.

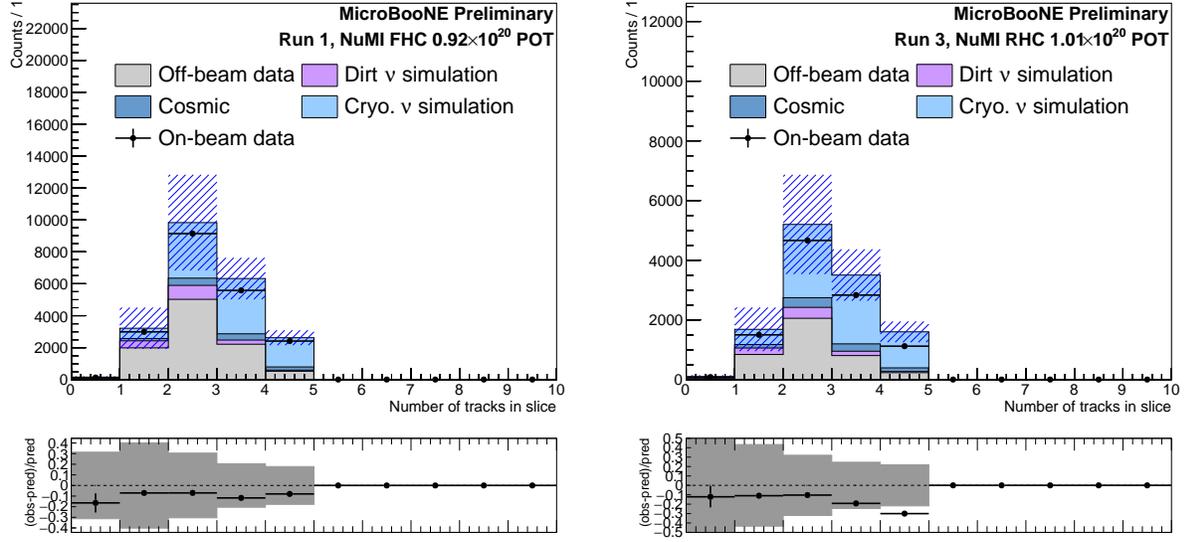


FIG. 14. Distribution showing the number of tracks in the slice for Run 1 (left) and Run 3 (right). The ‘cosmic’ background is cosmic contamination from the underlying event in simulated data. The hashed area is the systematic error band due to uncertainties in the flux and cross-section models, and simulation statistics. The ratio plot shows the total uncertainty in grey. A 10% (15%) deficit of data with respect to the model prediction in Run 1 (Run 3) is within the flux modelling uncertainty.

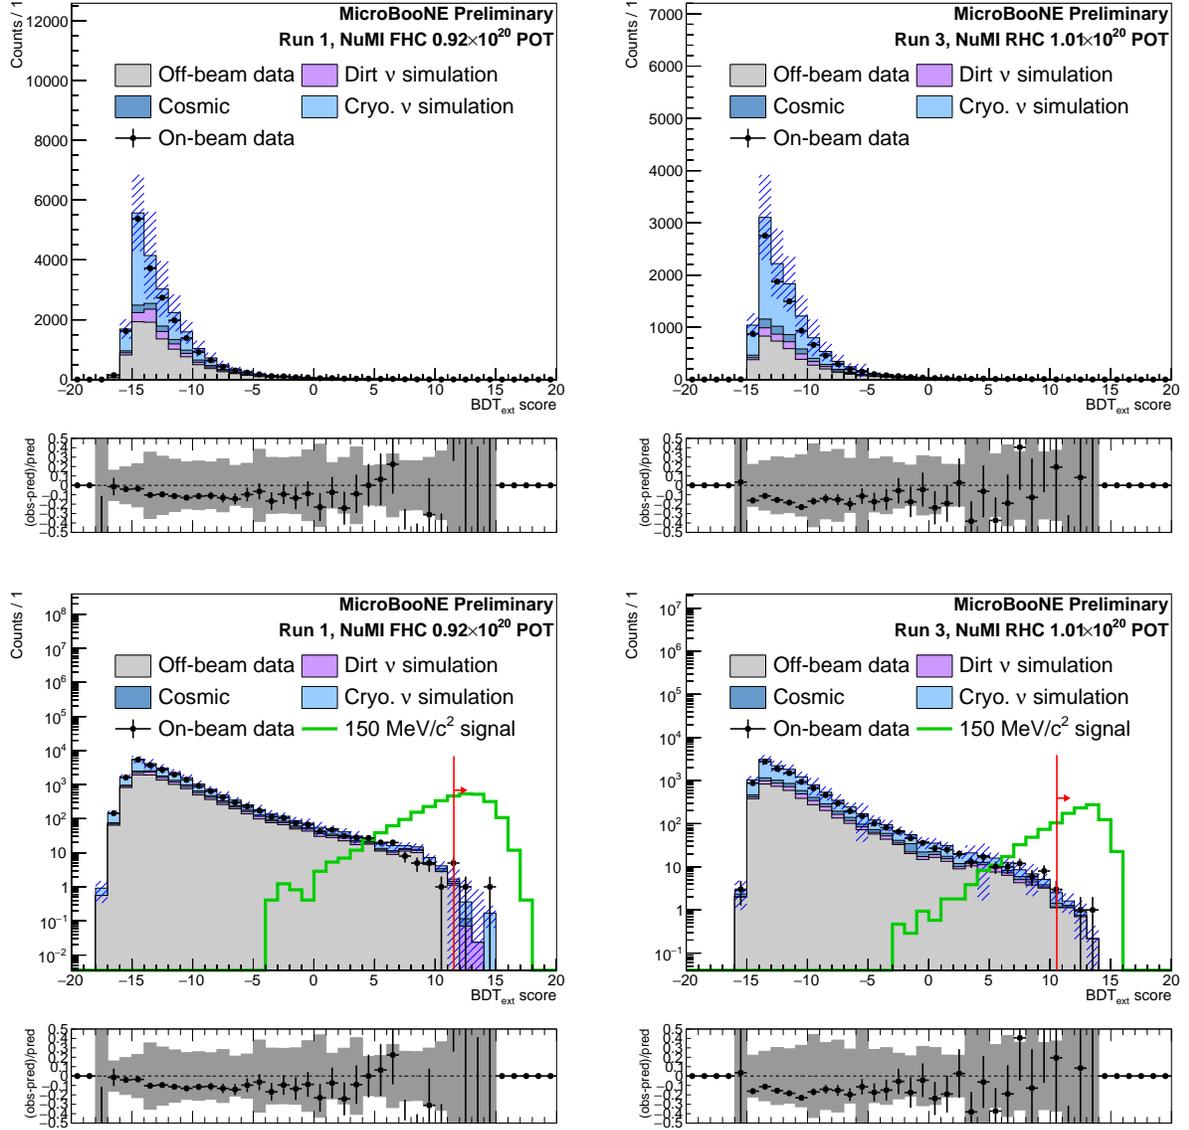


FIG. 15. Distribution showing the score of the cosmic-focused BDT for Run 1 (left) and Run 3 (right). The ‘cosmic’ background is cosmic contamination from the underlying event in simulated data. The hashed area is the systematic error band due to uncertainties in the flux and cross-section models, and simulation statistics. The ratio plot shows the total uncertainty in grey. A 10% (15%) deficit of data with respect to the model prediction in Run 1 (Run 3) is within the flux modelling uncertainty. The bottom panels show the same distributions but with a logarithmic scale. Also shown is the sensitivity-optimised cut value in red, and an example signal distribution in green for  $m_S = 150 \text{ MeV}/c^2$ , with arbitrary normalisation.

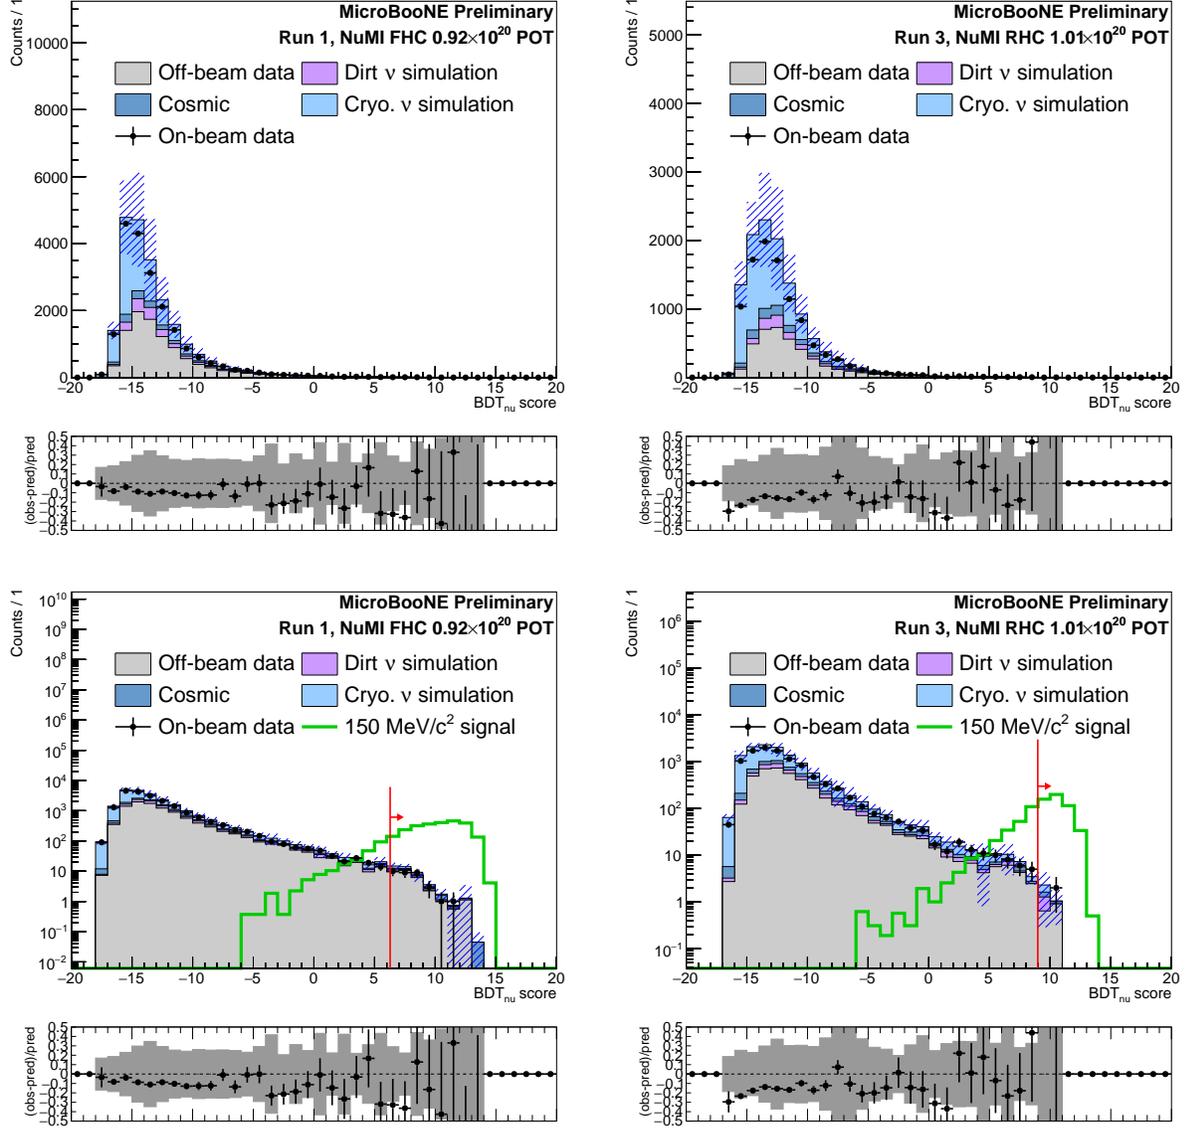


FIG. 16. Distribution showing the score of the neutrino-focused BDT for Run 1 (left) and Run 3 (right). The ‘cosmic’ background is cosmic contamination from the underlying event in simulated data. The hashed area is the systematic error band due to uncertainties in the flux and cross-section models, and simulation statistics. The ratio plot shows the total uncertainty in grey. A 10% (15%) deficit of data with respect to the model prediction in Run 1 (Run 3) is within the flux modelling uncertainty. The bottom panels show the same distributions but with a logarithmic scale. Also shown is the sensitivity-optimised cut value in red, and an example signal distribution in green for  $m_S = 150 \text{ MeV}/c^2$ , with arbitrary normalisation.

- 
- [1] B. Patt and F. Wilczek, Higgs-field portal into hidden sectors, (2006), arXiv:hep-ph/0605188.
- [2] B. Batell, J. Berger, and A. Ismail, Probing the Higgs Portal at the Fermilab Short-Baseline Neutrino Experiments, *Phys. Rev. D* **100**, 115039 (2019), arXiv:1909.11670 [hep-ph].
- [3] S. Shinohara (KOTO), Search for the rare decay at J-PARC KOTO experiment, *J. Phys. Conf. Ser.* **1526**, 012002 (2020).
- [4] D. Egana-Ugrinovic, S. Homiller, and P. Meade, Light Scalars and the KOTO Anomaly, *Phys. Rev. Lett.* **124**, 191801 (2020), arXiv:1911.10203 [hep-ph].
- [5] A. Artamonov *et al.* (BNL-E949), Study of the decay  $K^+ \rightarrow \pi^+ \nu \bar{\nu}$  in the momentum region  $140 < P_\pi < 199$  MeV/c, *Phys. Rev. D* **79**, 092004 (2009), arXiv:0903.0030 [hep-ex].
- [6] E. Cortina Gil *et al.* (NA62), An investigation of the very rare  $K^+ \rightarrow \pi^+ \nu \bar{\nu}$  decay, (2020), arXiv:2007.08218 [hep-ex].
- [7] F. Bergsma *et al.* (CHARM), Search for Axion Like Particle Production in 400-GeV Proton - Copper Interactions, *Phys. Lett. B* **157**, 458 (1985).
- [8] M. W. Winkler, Decay and detection of a light scalar boson mixing with the Higgs boson, *Phys. Rev. D* **99**, 015018 (2019), arXiv:1809.01876 [hep-ph].
- [9] P. B. Dev, R. N. Mohapatra, and Y. Zhang, Constraints on long-lived light scalars with flavor-changing couplings and the KOTO anomaly, *Phys. Rev. D* **101**, 075014 (2020), arXiv:1911.12334 [hep-ph].
- [10] R. Acciarri *et al.* (MicroBooNE), Design and Construction of the MicroBooNE Detector, *JINST* **12** (02), P02017, arXiv:1612.05824 [physics.ins-det].
- [11] C. Adams *et al.* (MicroBooNE), Design and construction of the MicroBooNE Cosmic Ray Tagger system, *JINST* **14** (04), P04004, arXiv:1901.02862 [physics.ins-det].
- [12] P. Adamson *et al.*, The NuMI Neutrino Beam, *Nucl. Instrum. Meth. A* **806**, 279 (2016), arXiv:1507.06690 [physics.acc-ph].
- [13] L. Aliaga *et al.* (MINERvA), Neutrino Flux Predictions for the NuMI Beam, *Phys. Rev. D* **94**, 092005 (2016), [Addendum: *Phys.Rev.D* **95**, 039903 (2017)], arXiv:1607.00704 [hep-ex].
- [14] S. Agostinelli *et al.* (GEANT4), GEANT4—a simulation toolkit, *Nucl. Instrum. Meth. A* **506**, 250 (2003).
- [15] A. Aguilar-Arevalo *et al.* (MiniBooNE), First Measurement of Monoenergetic Muon Neutrino Charged Current Interactions, *Phys. Rev. Lett.* **120**, 141802 (2018), arXiv:1801.03848 [hep-ex].
- [16] L. Aliaga Soplin, *Neutrino Flux Prediction for the NuMI Beamline*, Ph.D. thesis, William-Mary Coll. (2016).
- [17] C. Andreopoulos *et al.*, The GENIE Neutrino Monte Carlo Generator, *Nucl. Instrum. Meth. A* **614**, 87 (2010), arXiv:0905.2517 [hep-ph].
- [18] The MicroBooNE Collaboration, Neutrino interaction model and uncertainties for MicroBooNE analyses, MICROBOONE-NOTE-1074-PUB.
- [19] C. Adams *et al.* (MicroBooNE), Ionization electron signal processing in single phase LArTPCs. Part I. Algorithm Description and quantitative evaluation with MicroBooNE simulation, *JINST* **13** (07), P07006, arXiv:1802.08709 [physics.ins-det].
- [20] R. Acciarri *et al.* (MicroBooNE), The Pandora multi-algorithm approach to automated pattern recognition of cosmic-ray muon and neutrino events in the MicroBooNE detector, *Eur. Phys. J. C* **78**, 82 (2018), arXiv:1708.03135 [hep-ex].
- [21] T. Chen and C. Guestrin, XGBoost: A scalable tree boosting system, in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16 (ACM, New York, NY, USA, 2016) pp. 785–794.
- [22] C. Andreopoulos, C. Barry, S. Dytman, H. Gallagher, T. Golan, R. Hatcher, G. Perdue, and J. Yarba, The GENIE Neutrino Monte Carlo Generator: Physics and User Manual, (2015), arXiv:1510.05494 [hep-ph].
- [23] L. Moneta, K. Belasco, K. S. Cranmer, S. Kreiss, A. Lazzaro, D. Piparo, G. Schott, W. Verkerke, and M. Wolf, The RooStats Project, *PoS ACAT2010*, 057 (2010), arXiv:1009.1003 [physics.data-an].